Orbiting binary black hole evolutions with a multipatch high order finite-difference approach

Enrique Pazos,[1,2,3] Manuel Tiglio,[1,2] Matthew D. Duez,[4] Lawrence E. Kidder,[4] and Saul A. Teukolsky[4]

[1] *Center for Fundamental Physics, Department of Physics, University of Maryland, College Park, MD 20742, USA*
[2] *Center for Scientific Computation and Mathematical Modeling, University of Maryland, College Park, MD 20742, USA*
[3] *Departamento de Matemática, Universidad de San Carlos de Guatemala, Edificio T4, Facultad de Ingeniería, Ciudad Universitaria Z. 12, Guatemala*
[4] *Center for Radiophysics and Space Research, Cornell University, Ithaca, New York, 14853*

We present numerical simulations of orbiting black holes for around twelve cycles, using a high-order multipatch approach. Unlike some other approaches, the computational speed scales almost perfectly for thousands of processors. Multipatch methods are an alternative to AMR (adaptive mesh refinement), with benefits of simplicity and better scaling for improving the resolution in the wave zone. The results presented here pave the way for multipatch evolutions of black hole-neutron star and neutron star-neutron star binaries, where high resolution grids are needed to resolve details of the matter flow.

PACS numbers: 04.25.dk, 04.40.Dg, 04.30.Db, 95.30.Sf

arXiv:0904.0493v1 [gr-qc] 3 Apr 2009

I. INTRODUCTION

Mergers of binary compact objects (neutron stars or black holes) are expected to be the main sources of gravitational waves for the ground-based interferometric detectors LIGO, GEO, VIRGO, and TAMA. Neutron star-neutron star and black hole-neutron star binaries are also interesting because they are leading candidates for explaining the production of short-duration gamma-ray bursts and because gravitational wave signals from these events may encode information about the neutron star equation of state [1, 2, 3]. Such a merger can be accurately modeled only by the numerical evolution of the full Einstein field equations coupled (if a neutron star is present) to an evolution of the neutron star matter.

Because of advances in numerical relativity in recent years, stable evolutions can now be performed for most binary cases. Accuracy and speed are now the pressing numerical challenges: how to achieve the minimum error given limited time and computational resources. A good code should converge rapidly with increasing resolution to the exact solution. Its speed should scale well with the number of processors used in order to make good use of parallelization. Also, an efficient use of resources will require a grid well adapted to the problem at hand. This includes using a grid with the most appropriate shape. For example, it is reasonable to suppose that excision inner boundaries and outer boundaries should be spherical. A good grid will also use higher resolution where it is most needed. For example, although the grid must extend out into the wave zone to extract the gravitational wave signal, lower resolution is needed in the wave zone than is needed in the vicinity of a black hole or neutron star. The need for high resolution in neutron stars and black hole accretion disks can become particularly acute in cases of hydrodynamic or magnetohydrodynamic instabilities, such as convective, Kelvin-Helmholtz, or magnetorotational instabilities. In such cases, the length scale of the unstable modes can be much smaller than the radius of the star or disk, and the evolution will be qualitatively wrong if the instability is completely unresolved.

One technique that has been successfully used to deal with this problem is adaptive mesh refinement (AMR) [4, 5]. These AMR codes generally use overlapping Cartesian meshes of varying levels of refinement, with the finer meshes being used only where they are determined (by some algorithm) to be needed. In this paper, we present a different method of achieving efficient grid coverage, one that is algorithmically simpler and that possesses some unique advantages.

This different technique for evolving binary compact systems involves using multiple grid patches, each patch having its own shape, curvilinear coordinates and resolution. The basic ideas behind these multipatch methods have been worked out in earlier papers [6, 7, 8]. In these references some particular patch configurations using cubes and cubed-spheres were used. The cubed-sphere patches were used to construct grids with exactly spherical inner excision boundaries and outer boundaries. These methods are, hence, ideal for calculations that involve excision. (Using AMR with excision introduces a number of complications.) These techniques were then successfully used to simulate perturbed Kerr black holes [9, 10]. Multiple patches in cubed-sphere arrangement have also been used to evolve the shallow water equations [11] and to simulate hydrodynamic flows in black hole accretion disks [12, 13, 14].

Another multipatch approach has been used by the Cornell-Caltech group to evolve Einstein's equations for binary black hole [15] and black hole-neutron star [16] systems. In the binary black hole case derivatives in these simulations are computed pseudospectrally, rather than using finite differencing. While pseudospectral methods produce accurate results very efficiently for binary black

hole evolutions, they are much less cost effective for systems involving matter. One reason for this is that the discontinuities that naturally appear in the fluid flow at shocks and stellar surfaces destroy the exponential convergence of spectral methods. In fact, the Cornell-Caltech group found it necessary to evolve the fluid variables using finite differencing, while evolving the field variables pseudospectrally. This required two independent grids: the finite difference gridpoints used to evolve the fluid, and the collocation points of the pseudospectral code used to evolve the metric. For the two grids to communicate, variables had to be interpolated from one grid to the other each timestep, a process which consumed about one third of the CPU time in each simulation. Another problem with pseudospectral techniques is that they usually do not scale well to large numbers of processors. In regions without discontinuities, where spectral convergence is not lost, one cannot, for example, split one large domain into two domains with half the number of collocation points each without a significant loss in accuracy. On the other hand, accurate simulations of binary neutron star or black hole-neutron star mergers are not practical without many processors.

It would therefore seem preferable to evolve both the fluid and the metric with finite differencing. This could significantly improve the scalability, allowing simulations on hundreds or thousands of processors. It would also remove the need for two separate grids and the expensive interpolation between them. Multipatch techniques are the natural finite difference version of the Cornell-Caltech pseudospectral evolution algorithm. As a first step in that direction, in this paper we evolve a binary black hole system using multipatches together with high order finite-differencing operators. We show that our code converges rapidly, scales well to thousands of processors, and can stably simulate several orbits of the inspiral.

II. EVOLUTION EQUATIONS

At the continuum level, the techniques used in this paper are exactly those ones previously used by the Caltech-Cornell collaboration in binary black hole simulations. We use the first order form of the generalized harmonic system presented in [17]. The evolution variables in this formulation are the 4-metric g_{ab} and its first derivatives in space and time $\partial_c g_{ab}$. (The indices run from 0 to 3.) We use the constraint preserving boundary conditions of [17, 18, 19]. The evolution of the gauge is determined by the gauge source functions $H_a = -g^{cd}\Gamma_{acd}$, which are freely specifiable functions of space and time. In this paper, the gauge is set by choosing H_a to be constant in time in a coordinate system that comoves with the holes. This comoving coordinate system is determined using the same dual frame and control tracking mechanism as was used for the spectral binary black hole evolutions [20]. This technique uses two coordinate frames, which we label $x^{\bar{i}}$ and x^i. The coordinate frame $x^{\bar{i}}$ is set to be an asymptotically flat, inertial frame. All tensor components are evaluated with respect to this frame. The gridpoints are fixed in the computational frame x^i. By means of a mapping between the frames, the computational coordinates can be made to approximately comove with the system. For the runs in this paper, we track the binary using a simple combination of rotation and radial scaling:

$$\begin{aligned} \overline{t} &= t \\ \overline{x} &= a[x\cos(\theta) - y\sin(\theta)] \\ \overline{y} &= a[x\sin(\theta) + y\cos(\theta)] \\ \overline{z} &= az \ , \end{aligned} \tag{1}$$

where θ and a are functions of time which are evolved using a feedback mechanism to keep the location of the black holes fixed in the computational domain.

The differences between the simulations presented in this paper and the earlier spectral simulations the type of domain decomposition, and in the numerical techniques used to compute the right-hand sides of the evolution equations (e.g. how spatial derivatives are approximated). Our handling of these issues is described below.

III. INITIAL DATA

The initial data that we use here consists of a snapshot at a given time of the highest resolution 16-orbit simulation done by the Caltech-Cornell collaboration, which corresponds to the run 30c1 reported in Refs. [15, 21]. The starting time $t = 0$ in our simulations corresponds to the instant $t = 2887\,M$ of the 16-orbit simulation (with M being the sum of the irreducible masses of each black holes). From that point, the black holes orbit for about 6 orbits before merger, although our runs stop before the merger takes place.

This way of specifying the initial data has the advantage that there is no junk radiation present in the computational domain at our starting time. Since the domains and points used in this paper are different from those used in the spectral simulation, we spectrally interpolate the initial data to the multipatch domain.

The outer boundary of our domain is a sphere of radius $r = 144\,M$. This value is actually mapped to $r' = 105\,M$ by the dual-frame coordinate transformation, which scales and rotates the inertial coordinates into the comoving ones. The coordinate transformation is a simple rescaling of the radial coordinate $r' = a(t)r$ by a time dependent factor, and a rigid rotation about the z axis. Since the binary system has been evolving before our $t = 0$ time, the scale factor has a value $a = 0.727$ and the rotation angle is $\theta = 57.95$ radians at the beginning of our simulations. The black hole coordinate separation at the beginning of the 30c1 run is $14.44\,M$. At our time $t = 0$ the initial coordinate separation is $10.5\,M$.

IV. MULTI-BLOCK DOMAIN

A. Structure

We use two types of basic building patches to cover the whole computational domain. One is simply a cuboid with a Cartesian coordinate map. The other is a combination of six patches that we call a *juggling ball*. A juggling ball can assume two different configurations. The first of them is shown at the top of Fig. 1. It consists of a cube whose interior has been excised by a sphere. We will refer to it as an *inner* juggling ball because it is the one that we use to excise the interior of each black hole and to cover its immediate surroundings. The second configuration is shown at the bottom of Fig. 1 and consists of a sphere whose interior has been excised by a cube. We will call it an *outer* juggling ball because it is the one that covers the region away from the black holes, reaching to the outer boundary. Both types of juggling balls use a radial coordinate that adjusts smoothly to their geometry. Each surface of constant radial coordinate is endowed with six two-dimensional coordinate maps, in the same fashion as the cubed sphere [11]. In essence, the juggling ball is a collection of six patches, each of them topologically equivalent to a cube.[1]

The basic layout of the full domain used in this paper is shown in Fig. 2. The centers of the excised spheres (which will be inside each black hole) are located along the x axis at $x = \pm a$. Here we have used two inner juggling balls, one around each black hole. Their individual outer boundaries are cubes with sides of length $2a$. When they are put together, we end up with a cuboid of dimensions $4a \times 2a \times 2a$, with the longest side along the x axis. We surround this structure with six cuboid patches of dimensions $4a \times 2a \times a$, aligning them along the y and z axes. After doing so, we end up with a cubical domain with sides of length $4a$. To complete the patch system we add an outer juggling ball whose cubical interior holds the two inner juggling balls plus the six cuboids. The outer juggling ball enables us to shape the outer boundary into a sphere, in which case moving the boundary further out requires an increase in the number of grid points that scales as $O(N)$ (as opposed to $O(N^3)$).

The total number of patches in this basic configuration is 6 cuboidal patches + 6 × (3 juggling balls) = 24 patches.

None of the patches used in this paper overlap with any other (in which case they are usually called *blocks*). A given block communicates with adjacent ones only by the two-dimensional common surface between them. Accordingly, we handle parallelization by assigning one block per processor, in this way minimizing communication between processors.

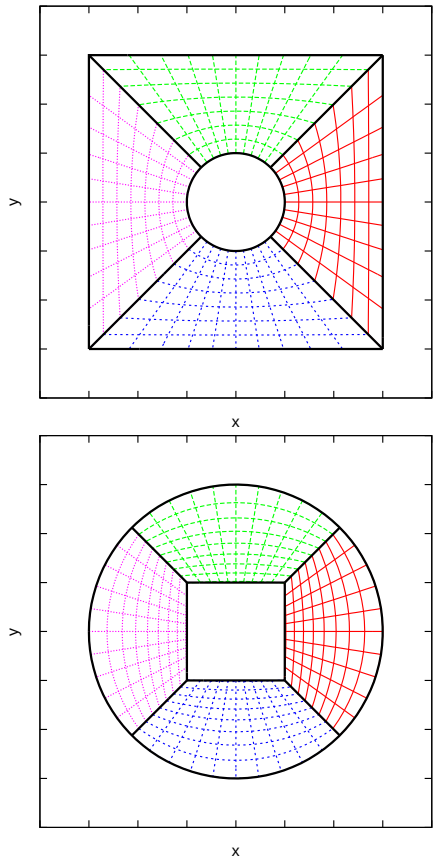

Figure 1: Equatorial cross-section of an inner juggling ball (top). Black lines denote the block boundaries. Colored lines represent the coordinate grid of each block. Equatorial cross-section of an outer juggling ball (bottom).

In this basic 24-block domain case, we would use exactly 24 processors, which is a fairly small number for a binary black hole simulation. In order to achieve higher resolutions by increasing the amount of points per block, we subdivide the existing blocks into smaller pieces. Since the topology of each block is cubical, it is straightforward to subdivide them. The guiding principle that we use to accomplish the subdivision is to keep the same number of points per block for every single block. Although this condition is not necessary, it is convenient because it balances the computational load across all the processors.

For the runs presented here, we used 192- and 384-block domains. The first case is obtained by subdividing the inner juggling balls uniformly in the radial direction 7 times. The 6 cuboids are split by a factor of 2 and the outer juggling ball is divided 4 times in the radial and twice in each transverse direction. The 384-block case is derived from the 192-block one by further split of each block in the radial direction by a factor of 2.

Figure 3 shows the multipatch structures used in this paper.

[1] The name juggling ball was chosen because some real juggling balls have a set of six quadrilateral-shaped designs on their surface.

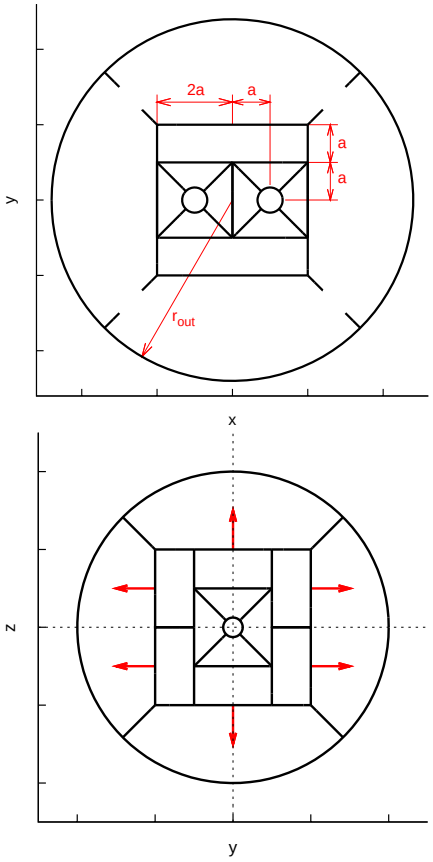

Figure 2: Equatorial cut of the computational domain (top). Schematic figure showing the direction considered as radial (red arrows) for the cuboidal blocks (bottom).

B. Numerical techniques

In our simulations we use the D_{8-4} summation-by-parts (SBP) finite difference operator and its associated dissipation constructed in [7]. The naming convention is meant to indicate that the derivative is 8th order accurate in the bulk of each block but only 4th order accurate near inter-block boundaries. The derivative in the interior of each block is a centered one and is modified near boundaries so as to satisfy the SBP property with respect to a diagonal norm; this is the cause of the drop in convergence. Information across sub-domains is communicated using characteristic variables and a penalty method (see [6, 7, 8] for more details).

The combination of these techniques guarantees numerical stability, but at the expense of the drop in convergence order near boundaries. For example, in the D_{8-4} case there are eight points near *each* boundary where the scheme is fourth order. For technical reasons explained below, in the simulations of this paper we use a rather large number of blocks and processors (192 and 384), with a very small load on each. As a consequence, the scheme is fourth order nearly everywhere and we expect our simulations to be 4th order convergent. This is indeed what our simulations below show.

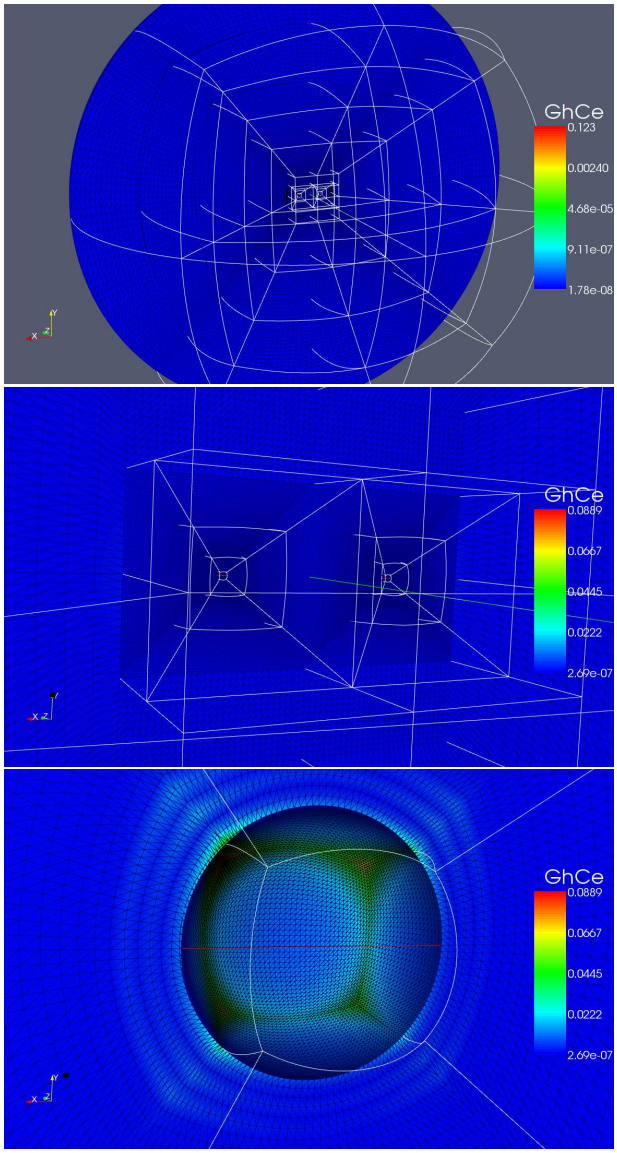

Figure 3: Computational domain used in the simulations of this paper

C. Resolution

One of the features that a multipatch method offers is the flexibility to increase only the radial resolution while keeping the angular resolution constant. Given that the angular profile of the waveforms is dominated by a few low-ℓ modes, once a sufficient angular resolution is used the truncation error will be dominated by the radial resolution.

The approximate spherical symmetry in the vicinity of each black hole and at large distances from them allows the radial direction to be naturally defined for each juggling ball block. However, for the cuboidal blocks there

is some arbitrariness in how to choose the radial direction. In practice, a radial direction for these blocks is useful only to define the direction along which resolution will be increased. In Fig. 2 the radial directions for the cuboidal blocks are indicated with arrows.

We use an angular resolution of $\pi/58$ around each black hole and twice as much in the outer blocks. That is, there are 116 points along an equatorial line around each black hole and twice that number in the distant wave region. This is somehow inefficient since the solution is over-resolved in the angular directions compared to the radial one (especially in the wave region). The motivation behind this choice was to allow the grid points at the boundary faces of adjacent blocks to be in one-to-one correspondence with each other. In this way the communication of the characteristic modes at the inter-patch boundaries does not require interpolation.

In Table I we show the total number of points in the whole domain and per block for the simulations of this paper. We increase resolution only along the radial direction, by the same number of points in all the domains. In our setup all blocks have the same number of points. Since parallelization is handled by assigning one block per processor, this guarantees a homogeneous load distribution.

The number of points shown in Table I is actually not large for a fully three-dimensional (i.e. no symmetries imposed) finite-difference simulation. For example, we can compare these numbers to a binary black hole evolution with around the same number of orbits using Cartesian grids and adaptive mesh refinement [22]. A typical state-of-the-art simulation uses six refinement levels around each black hole with 64^3 points on each level, and four coarse grid levels with 128^3 points. This amounts to a total of $2\times 6\times 64^3 + 4\times 128^3 \approx 226^3$ points. In the case of non-spinning, equal-mass black holes one can make use of the symmetry of the problem and reduce the total number of points to $6\times 64^3 + 128^3 \approx 154^3$.

We have tested the performance of our multipatch parallelization scheme for the evolution of a single black hole. In Fig. 4 we show a strong scaling test for up to $3,000$ processors (cores), in which the total number of points is kept fixed while increasing the number of processors. We see that the speed of the code has a linear dependence on the number of processors. Similarly, in Fig. 5 we show a weak scaling test, where the load per processor is kept fixed while increasing the number of processors used. The drop in speed in this case is about 15% over a range of 10 to 3,000 processors. We have not attempted to go beyond this number of cores.

The phase errors in the waveforms shown in the next section are rather large compared to state-of-the-art simulations (in particular, compared to an AMR one such as the one mentioned above). Since the code scales well and the number of points used in this paper (shown in Table I) is reasonable for a finite-difference evolution, in principle we could improve the accuracy of the simulations shown in the next section while still using modest computational resources. What has prevented us from doing so is a purely technical obstacle. The computational infrastructure used in this paper, SpEC, was originally designed for pseudo-spectral evolutions, which are extremely efficient in terms of memory. For that reason SpEC currently stores in memory many more variables than are actually needed for evolving the system. As a result, in our FD simulations because of memory constraints we actually end up using a few cores per node and a rather large number of nodes. We plan to improve SpEC's use of memory soon to eliminate this limitation. However, for the demonstrations in this paper, the current resolutions are sufficient.

$N_r \times N_{\rm ang} \times N_{\rm blocks} = N_{\rm total}$	speed ($\rm h^{-1}$)	CPU (h)
$19 \times 29^2 \times 192 = 145^3$	2.83	67844
$22 \times 29^2 \times 192 = 153^3$	1.86	103226
$16 \times 29^2 \times 384 = 173^3$	2.42	158678

Table I: Speed and CPU time for three resolutions. N_r and $N_{\rm ang}$ are the radial and angular number of points per block, respectively, as described in the text. The speed is expressed in units of the total irreducible mass per hour.

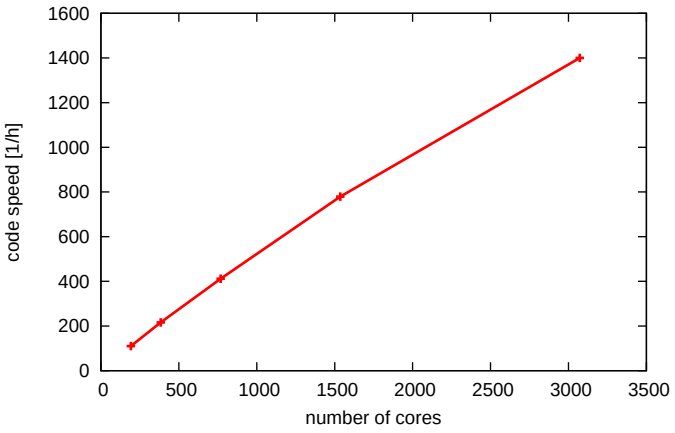

Figure 4: Strong scaling test for a single black hole. The speed of the code depends essentially linearly on the number of processors, almost perfect scaling.

V. RESULTS

Figure 6 shows the location of the centroids of the apparent horizons for the highest resolution simulation. The black holes complete about six orbits before reaching the merger regime.

A. Convergence of the constraints

A way of checking the consistency of the numerical solution is monitoring the constraint violations, since they

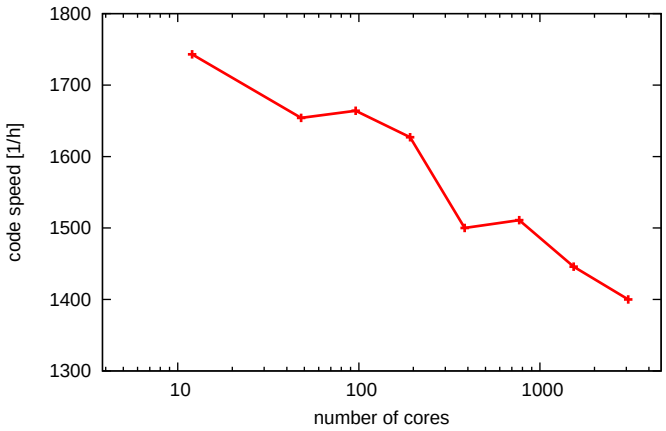

Figure 5: Weak scaling test for a single black hole. There is only a 15% drop in speed as the number of processors is increased while keeping the load per processor fixed.

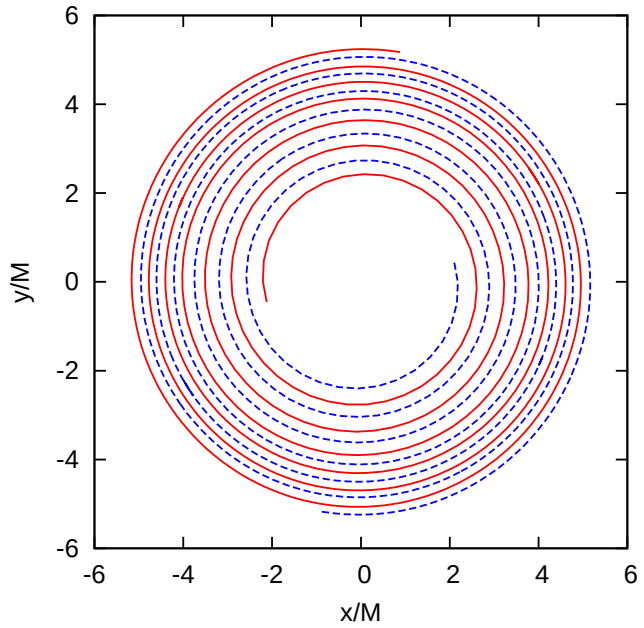

Figure 6: Black hole orbits.

are not enforced during the evolution. In Fig. 7 we plot the L^2 norm of all the constraint fields of the first order generalized harmonic system, normalized by the L^2 norm of the spatial gradients of the dynamical fields, as defined in [17]. We show three runs with different resolutions.

Figure 8 shows the convergence exponent of the L^2 norm of the normalized constraint violations, which is around four, as expected (cf. Sec. IV B). The convergence exponent n is defined as

$$\frac{\beta^n - 1}{\beta^n - \alpha^n} = \frac{C_1 - C_3}{C_2 - C_3} \qquad (2)$$

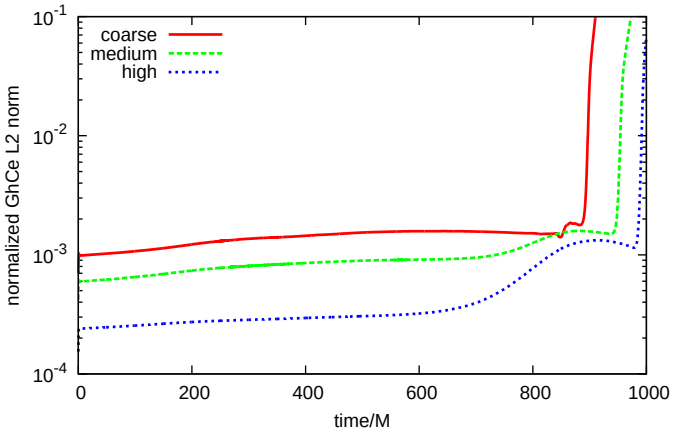

Figure 7: L^2 norm of the normalized constraints.

where α is the ratio between the medium and coarse resolution and β, the ratio between the fine and coarse one. C_1, C_2, and C_3 represent a given quantity at coarse, medium and fine resolutions, respectively.

The uniform convergence is lost around $t \sim 800\,M$, at which time the values for the coarse and medium resolutions intersect, as is seen in Fig. 7.

We stop our simulations when the characteristic speeds at the excision boundary change sign, which means that there is spurious information entering the domain. That moment is characterized by a blow-up of the constraints. This feature is due to the inadequacy of the rather simple gauge conditions used here at times close to merger. At the time the simulations of this paper were performed, we used the same simple conditions used then by the Caltech-Cornell collaboration, namely, keeping the gauge source functions fixed in the comoving frame. Since then, better conditions have been developed, which do allow simulations to go through merger and ringdown [15]. For the purposes of this paper, however, following six orbits of an inspiral is sufficient.

B. Waveforms

Waveforms are computed via the Newman-Penrose curvature scalar Ψ_4 as in [23]. Subsequently we decompose Ψ_4 in spin-weighted spherical harmonics ${}_{-2}Y_{\ell m}(\theta,\phi)$. We focus our discussion to the $\ell = 2$, $m = 2$ mode. The extraction is done at $r = 50\,M$.

Figure 9 shows the real component of Ψ_4

We see that they all agree at early times and drift apart during the later stages of the evolution. A more meaningful comparison is shown in Figs. 10 and 11, where we plot amplitude and phase of the extracted wave. The differences between the finite differences waveforms and the spectral one are shown in Figs. 12 and 13 for the amplitude and phase, respectively.

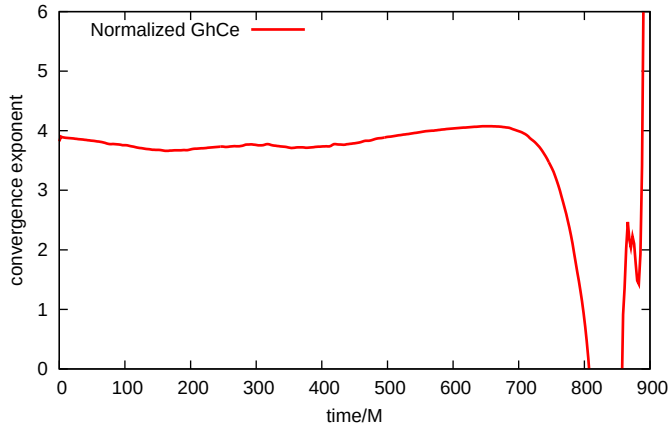

Figure 8: Convergence exponent for the L^2 norm of the normalized constraints.

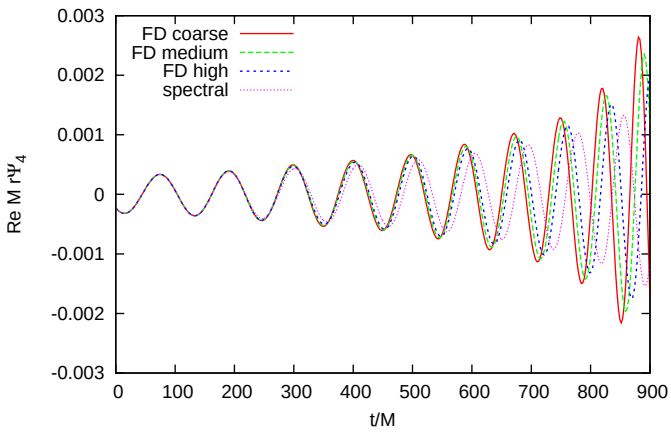

Figure 9: Real part of Ψ_4 for the finite difference and spectral results.

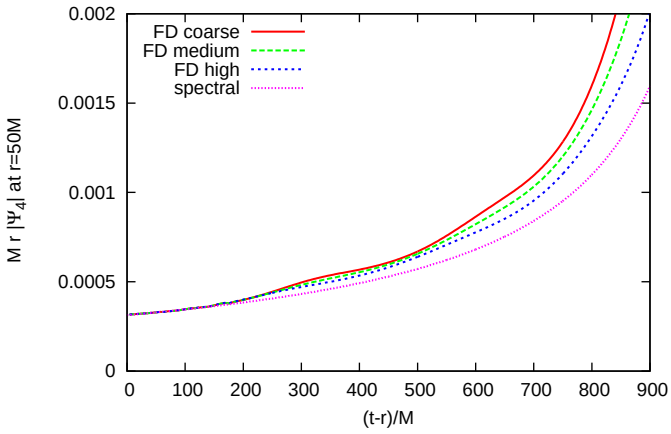

Figure 10: Ψ_4 amplitude for the finite difference and spectral results.

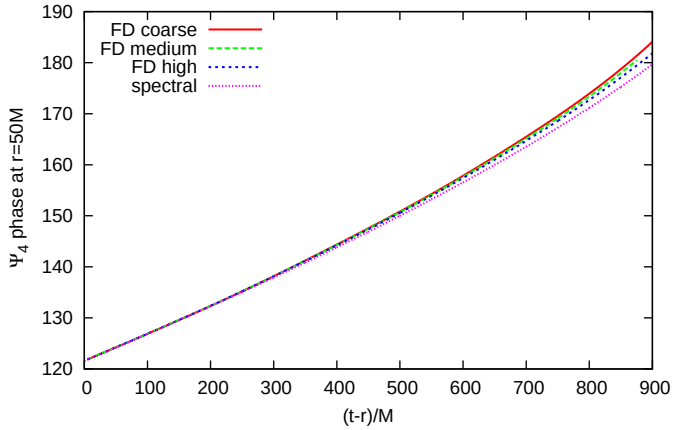

Figure 11: Ψ_4 phase for the finite difference and spectral results.

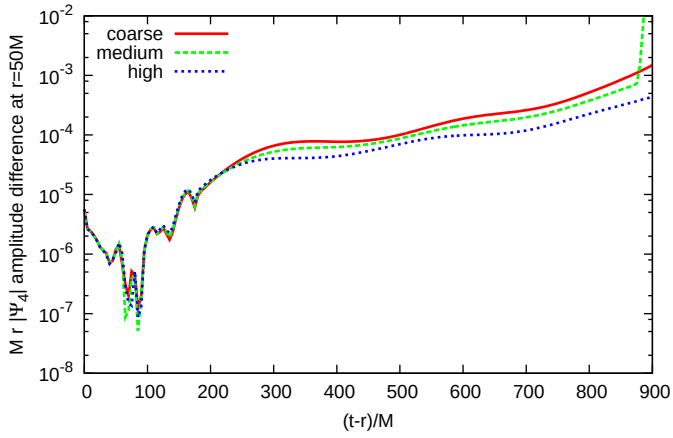

Figure 12: Differences in the Ψ_4 amplitude between the finite difference and the spectral results.

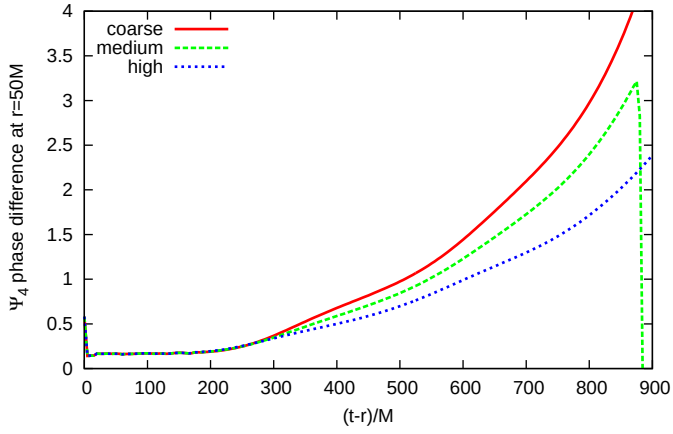

Figure 13: Differences in the Ψ_4 phase between the finite difference and the spectral results.

VI. REMARKS

In this paper we have shown that we can evolve orbiting black holes in a stable way using a high-order multipatch approach and that this method scales well with the number of processors. As a result, we expect to be able to achieve good accuracy while still using only modest computational resources. These results also suggest that multipatch methods are an excellent alternative to AMR, with benefits of simplicity and $O(N)$ scaling for improving resolution in the wave zone. Finally, multipatch methods will allow one to use the same grid to evolve both metric and matter fields for a binary pair composed of a black hole and neutron star, allowing the advantages of high-order methods without the drawbacks of a hybrid spectral-finite difference approach.

Acknowledgments

The numerical simulations presented here were performed using the Spectral Einstein Code (SpEC) developed at Caltech and Cornell primarily by Larry Kidder, Mark Scheel and Harald Pfeiffer.

This research has been supported in part by NSF Grant No. PHY-0801213 to the University of Maryland, by NSF Grants No. PHY-0652952, No. DMS-0553677, and No. PHY-0652929 and by a grant from the Sherman Fairchild Foundation to Cornell, by the TeraGrid allocation TG-MCA02N014 to Louisiana State University, and by supercomputing resources at LONI.

We thank Oleg Korobkin for his help on the early stages of this project.

[1] R. Oechslin and H. T. Janka, Phys. Rev. Lett. **99**, 121102 (2007), astro-ph/0702228.
[2] J. S. Read et al. (2009), 0901.3258.
[3] M. Vallisneri, Phys. Rev. Lett. **84**, 3519 (2000), gr-qc/9912026.
[4] T. Yamamoto, M. Shibata, and K. Taniguchi, Phys. Rev. **D78**, 064054 (2008), 0806.4007.
[5] L. Baiotti, B. Giacomazzo, and L. Rezzolla, Phys. Rev. **D78**, 084033 (2008), 0804.0594.
[6] L. Lehner, O. Reula, and M. Tiglio, Class. Quant. Grav. **22**, 5283 (2005), gr-qc/0507004.
[7] P. Diener, E. N. Dorband, E. Schnetter, and M. Tiglio, J. Sci. Comput. **32**, 109 (2007), gr-qc/0512001.
[8] E. Schnetter, P. Diener, E. N. Dorband, and M. Tiglio, Class. Quant. Grav. **23**, S553 (2006), gr-qc/0602104.
[9] E. N. Dorband, E. Berti, P. Diener, E. Schnetter, and M. Tiglio, Phys. Rev. **D74**, 084028 (2006), gr-qc/0608091.
[10] E. Pazos, E. N. Dorband, A. Nagar, C. Palenzuela, E. Schnetter, and M. Tiglio, Class. Quantum Grav. **24**, S341 (2007), gr-qc/0612149.
[11] C. Ronchi, R. Iacono, and P. S. Paolucci, J. Comp. Phys **124**, 93 (1996).
[12] A. V. Koldoba, M. M. Romanova, G. V. Ustyugova, and R. V. E. Lovelace, Astrophys. J. **576**, L53 (2002), astro-ph/0209598.
[13] B. Zink, E. Schnetter, and M. Tiglio, Phys. Rev. **D77**, 103015 (2008), 0712.0353.
[14] P. C. Fragile, C. C. Lindner, P. Anninos, and J. D. Salmonson, Astrophys. J. **691**, 482 (2009), 0809.3819.
[15] M. A. Scheel et al., Phys. Rev. **D79**, 024003 (2009), 0810.1767.
[16] M. D. Duez et al., Phys. Rev. **D78**, 104015 (2008), 0809.0002.
[17] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, Class. Quant. Grav. **23**, S447 (2006), gr-qc/0512093.
[18] O. Rinne, Class. Quant. Grav. **23**, 6275 (2006), gr-qc/0606053.
[19] O. Rinne, L. Lindblom, and M. A. Scheel, Class. Quant. Grav. **24**, 4053 (2007), 0704.0782.
[20] M. A. Scheel et al., Phys. Rev. **D74**, 104006 (2006), gr-qc/0607056.
[21] M. Boyle et al., Phys. Rev. **D76**, 124038 (2007), 0710.0158.
[22] F. Herrmann (2009), private communication.
[23] H. P. Pfeiffer et al., Class. Quant. Grav. **24**, S59 (2007), gr-qc/0702106.